

General Complex Polynomial Root Solver and Its Further Optimization for Binary Microlenses

Jan Skowron & Andrew Gould

Department of Astronomy, Ohio State University, 140 W. 18th Ave., Columbus, OH 43210, USA; jskowron,gould@astronomy.ohio-state.edu

ABSTRACT

We present a new algorithm to solve polynomial equations, and publish its code, which is 1.6–3 times faster than the `ZROOTS` subroutine that is commercially available from *Numerical Recipes*, depending on application. The largest improvement, when compared to naive solvers, comes from a fail-safe procedure that permits us to skip the majority of the calculations in the great majority of cases, without risking catastrophic failure in the few cases that these are actually required. Second, we identify a discriminant that enables a rational choice between Laguerre’s Method and Newton’s Method (or a new intermediate method) on a case-by-case basis. We briefly review the history of root solving and demonstrate that “Newton’s Method” was discovered neither by Newton (1671) nor by Raphson (1690), but only by Simpson (1740). Some of the arguments leading to this conclusion were first given by the British historian of science Nick Kollerstrom in 1992, but these do not appear to have penetrated the astronomical community. Finally, we argue that *Numerical Recipes* should voluntarily surrender its copyright protection for non-profit applications, despite the fact that, in this particular case, such protection was the major stimulant for developing our improved algorithm.

Subject headings:– gravitational lensing – planetary systems – methods numerical

1. Introduction

Numerical methods designed for finding zeros of a function were discovered hundreds of years ago. The first such processes were found by Newton (1671) and Halley (1694). These

algorithms were improved and made more universal by 18th and 19th century mathematicians, notably Laguerre (1880) whose work we repeatedly use in this paper.

The binary lens equation gives rise to a fifth-order complex polynomial (Witt & Mao 1995), which must be solved by numerical methods (Abel 1824, 1826). Solving this equation is central to nearly all methods of modeling observed light curves due to binary lenses, including those with a low-mass, i.e., planetary companion.

Because of the elementary nature and long history of the problem of solving polynomials, it is generally believed that further optimization of root solvers is not possible. For example Bozza (2010), when discussing his algorithm for calculating binary-lens magnifications, writes: “We find that roughly 80 per cent of the machine time is spent in the root finding routine, for which there is basically no hope of further optimization.”

The value of 80%, cited above, refers to the calculation by contour integration (Dominik 1993, 1998; Gould & Gauchere 1997), one of the most versatile and commonly used methods in microlensing. Increasing the speed of a root finding algorithm by a factor of 2 would decrease the time of execution by a factor 1.6 and thus, effectively, decrease the need of new computers, saving money, electricity and other resources. This is a task worth pursuing.

Section 2 of the paper presents a new algorithm for solving general complex polynomial equations. Like other such algorithms, it locates roots one by one numerically, and then divides the original polynomial by the found root before searching for the next one. However, in contrast to previous algorithms, this new one can, at each step, efficiently decide whether to use Laguerre’s method, or whether to choose a faster root finding method, either Newton’s method (Simpson 1740) or a new intermediate method, which is presented in Section 2.2.

Section 3 is focused on binary microlensing. We find that some significant improvements are possible by making use of specific features of the binary lens equation that do not apply to the full range of complex polynomials for which this more general algorithm was developed. We discuss these further optimizations and describe the codes there.

Moreover, in the course of implementing these improvements, we discovered several important results on the limits of precision of the lens solver. It is possible that these are already “known” by numerical mathematicians, but do not appear to be known by astronomers. We were unable to determine whether these results are known because we find the numerical literature to be virtually impenetrable. We expect that other astronomers may suffer similar difficulties. We discuss these results in Section 5.

In Appendix A we review the study that led us to identify Thomas Simpson as the discoverer of the so-called “Newton’s Method”. Appendix B contains discussion of *Numerical*

Recipes (Press et al. 1992, NR) code copyright protection and suggestion to waive it for non-profit and academic uses.

All codes described in this paper are open-source, are provided on the author’s web page¹ and are attached to the arxiv.org version of the paper. Please cite this paper if you are using these codes for a scientific work². A list of subroutines can be found in Appendix C.

2. New algorithm for solving complex polynomials

Laguerre’s Method (Laguerre 1880) is an iterative method similar to Newton’s Method (Simpson 1740) but is more stable because it uses first and second derivative. In addition, it converges faster, although each step takes more time to calculate. Non-convergent cycles are very rare and can be broken by introducing a simple modification to the method that does not hamper the overall speed of the algorithm.

To robustly find all n roots of the n th order polynomial, making use of a single-root finding numerical method, such as Laguerre’s method, the algorithm should search for roots successively one by one. After every single root is found, it should be eliminated from the polynomial by division, so the next searches yield a different root. When the degree of the resulting polynomial reaches zero, all roots are found. We call this step “robust”. Each search can be started from point (0,0) or from other initial guess on the positions of the root.

The division process introduces numerical noise. Hence, once the roots are all located in this fashion, each one can be “polished” by applying the same method used in “robust” to the full polynomial in the neighborhood of initially-located root.

The subroutine `CMPLX_ROOTS_GEN` that we introduce in this section, contains a general polynomial solver that has the above overall structure, but employs a new algorithm for single root searches (§ 2.2). We also discuss optimized implementation of Laguerre’s method (§ 2.1), a stopping criterion for polynomials (§ 2.4), and the detailed structure of the whole algorithm (§ 2.5).

Our algorithm is implemented in double precision, with all constants being in sync with this level of precision.

¹http://www.astrouw.edu.pl/~jskowron/cmplx_roots_sg/

²Open-source license’s details are described on the mentioned web page and inside the code package in the LICENSE file.

2.1. Implementation of Laguerre’s Method

Laguerre’s Method is implemented in subroutine `CMPLX_LAGUERRE` using the formula

$$\Delta_{\text{Lag}} = \Delta_{\text{Newt}} \left(\frac{1}{n} + \frac{n-1}{n} \sqrt{1 - \frac{n}{n-1} F(z)} \right)^{-1} \quad (\text{for } n > 1), \quad (1)$$

where

$$\Delta_{\text{Newt}} \equiv -\frac{p(z)}{p'(z)}; \quad F(z) \equiv -\frac{p''(z)}{p'(z)} \Delta_{\text{Newt}}. \quad (2)$$

and $p(z)$ is an n th order polynomial. To avoid division by small numbers, we choose the value of the square root with positive real component. Most compilers return this automatically, but we leave a simple check of the sign in the code to make it suitable for every compiler. We encourage users to test what convention is used within their compilers and to remove this check if appropriate.

In each loop of the algorithm we calculate the value of the polynomial, as well as its first two derivatives. We check the stopping criterion (Adams 1967, see Section 2.4) immediately after the value of the polynomial is evaluated. If the value of the polynomial is one order of magnitude lower than the stopping criterion for a given point, we return immediately. If the value meets the stopping criterion but there is still room for improvement (see Section 2.4), we set a flag that makes the subroutine return right after the next iteration is completed, without need for recalculating the value of the polynomial at that point.

This implementation of Laguerre’s method is faster than the one in `ZROOTS` from NR. The biggest improvement (by $\sim 10\%$) comes from the fact that to decide which sign of the square root should be taken we do not evaluate both denominators and do not calculate their absolute values, whereas `ZROOTS` uses a different representation of Δ_{Lag}

$$\Delta_{\text{Lag}} = n \left(G \pm \sqrt{(n-1)(nH - G^2)} \right)^{-1} \quad (3)$$

where

$$G(z) = \frac{p'(z)}{p(z)}; \quad H(z) = G(z)^2 - \frac{p''(z)}{p(z)}. \quad (4)$$

and does undertake these steps. We note, however, that even in this form one could substitute calculation of absolute values by a simpler check, which would recover most of the execution-time difference. That is, one could choose the “+” root if $\text{Re}(\overline{G} \sqrt{(n-1)(nH - G^2)}) > 0$.

To avoid cycles in the Laguerre’s algorithm, every 10th step we modify Δ_{Lag} by multiplying it by a number between 0.0 and 1.0.

Our implementation of Laguerre’s method (`CMPLX_LAGUERRE`) calculates the stopping criterion in every loop, since it is used to locate a root on a broad plane, over which the value

of the stopping criterion can change considerably. This is in contrast to our implementation of Newton’s method (`CMLX_NEWTON_SPEC`), which we use mainly for “root polishing”. Thus we calculate the stopping criterion on the first iteration of Newton’s method, and then only every 10th step (as a failsafe procedure). This saves significant time, since evaluating the stopping criterion takes about 1/3 of the time of a single Newton’s step.

The suffix `_SPEC` in the name of the subroutine is meant to suggest that this routine is not a generic implementations of the Newton’s method, but rather is specifically tailored for certain tasks in the broader algorithm. Users should not reuse provided routines without broader understanding of the choices made and implemented within them.

2.2. New algorithm

In traditional root solvers, Laguerre iterations consume the great majority of the computation time. For the binary-lens problem (5th order polynomial), four Laguerre searches are required to find the roots and five more to polish them. A close study of this algorithm is therefore worthwhile to determine under what conditions it can be simplified or replaced by Newton’s Method.

For an n th order polynomial $p(z)$, the Laguerre iteration, Δ_{Lag} , can be written in the form of Equation (1). This form has several interesting features. Most notably, it shows immediately that Laguerre iterations require about twice as much time as Newton iterations. First, once Δ_{Newt} is calculated, it requires only slightly fewer additional operations (so slightly less additional time) to calculate $F(z)$. Second, substantial calculations are required after $F(z)$ is found. Finally, Equation (1) actually has two roots, and a check may be required to see which leads to the larger modulus of the denominator, in order to avoid dividing by small numbers.

When $F(z)$ is small, we can Taylor expand Δ_{Lag} to obtain:

$$\Delta_{\text{Lag}} \approx \Delta_{\text{Newt}} \left(1 + \frac{F(z)}{2} + \frac{3}{8} \frac{n}{(n-1)} F(z)^2 + \dots \right). \quad (5)$$

In particular, at the limit of vanishing F this is exactly equal to Newton’s Method. In the neighborhood of an isolated root, $F(z)$ is of the same order as Δ_{Newt} . This immediately implies that for polishing isolated roots, one can dispense with Laguerre’s Method altogether and just use Newton’s Method. This simplification is actually of central importance from a practical standpoint in microlensing because the lens-solver is the rate-limiting step for the contour method (not ray-shooting), and the overwhelming majority of contour-method calls to the lens-solver should use only polishing.

Now, there is some cost to this because Laguerre’s Method does converge in fewer steps than Newton’s Method, but we find by experimentation that the number of steps is $\sim 30\%$ fewer, while the computation time of each step is about 2 times longer.

Moreover, Equation (5) suggests a new algorithm. That is, the value of $F(z)$ can be used as a discriminant to decide which method to use. If $|F|$ is very small, Newton’s Method can be used successfully; while for large values of $|F|$, the full Laguerre Method will be more suitable. In between these two regimes, an intermediate method can be used, i.e., the second order Taylor expansion:

$$\Delta_{\text{SG}} = \Delta_{\text{Newt}} \left(1 + \frac{F(z)}{2} \right). \quad (6)$$

Computation of Δ_{SG} is fast because F has already been evaluated, and it has the advantages of skipping the square-root evaluation, skipping one division, and determining the sign of the square root.

Laguerre’s Method is a second order method suitable for polynomials. We now show that the intermediate method presented in the Equation (6) is in fact a Second-order General method for all differentiable functions.

2.2.1. General 2nd order method

By Taylor expanding the general function $f(z)$ in the neighborhood of its root (z_0) we have:

$$0 = f(z_0) = f(z) + f'(z)(z_0 - z) + \frac{f''(z)}{2}(z_0 - z)^2 + \frac{f'''(z)}{6}(z_0 - z)^3 \dots \quad (7)$$

The first 3 terms of the expansion constitute a quadratic equation in the variable $(z_0 - z)$. One of the roots of this equation is approximated by:

$$z_0 - z = -\frac{f(z)}{f'(z)} \left(1 + \frac{1}{2} \frac{f(z)f''(z)}{f'(z)^2} \right) = \Delta_{\text{Newt}} \left(1 + \frac{F(z)}{2} \right) = \Delta_{\text{SG}}, \quad (8)$$

where

$$\Delta_{\text{Newt}} \equiv -\frac{f(z)}{f'(z)}; \quad F(z) \equiv \frac{f(z)f''(z)}{f'(z)^2}, \quad (9)$$

provided that F can be regarded as small. This new Second-order General method is similar to Halley’s Method (Halley 1694) developed for solving polynomials – however, its origin is clearer.

Since F needs to be calculated in each step to get new approximation of the root, one could leave in place the next term in F in the expansion of the root of the quadratic:

$$z_0 - z = \Delta_{\text{Newt}} \left(1 + \frac{F(z)}{2} + \frac{F(z)^2}{2} \right). \quad (10)$$

We will see immediately below, that this is equivalent to the 3rd order method, in the case that the third derivative of the function is negligible ($E(z) \ll 1$, see Eq. (11) for definition).

2.2.2. General 3rd order method

We now derive a third order method that uses third derivative of the function ($f'''(z)$). Inspired by the form of the discriminative function F and by a variation of Householder's 3rd order method, we construct the function:

$$E(z) \equiv \frac{f(z)^2 f'''(z)}{f'(z)^3}, \quad (11)$$

We note that the use of the function E instead of f''' and the use of $q = (z_0 - z)/\Delta_{\text{Newt}}$ as a variable, allow us to write Equation (7), up to the 3rd order, in the simple form:

$$0 = -1 + q - \frac{F}{2}q^2 + \frac{E}{6}q^3. \quad (12)$$

One of the solutions to this equation is approximated by:

$$z_0 - z = \Delta_{\text{Newt}} \left(1 + \frac{F(z)}{2} + \frac{F(z)^2}{2} - \frac{E(z)}{6} - \frac{5}{12}F(z)E(z) + \frac{E(z)^2}{12} + \dots \right), \quad (13)$$

which gives us a new method of finding roots complete to third order in Δ_{Newt} , $F(z)$ and $E(z)$.

2.2.3. $F(z)$: Dynamic discriminant of iteration method

We note that $F(z)$ must be calculated as an intermediate step in Laguerre's Method. However, before actually implementing Laguerre, we first evaluate³ $|F(z)|$. If $0.05 < |F(z)| < 0.5$, we use the intermediate method, wherein we approximate $\Delta = \Delta_{\text{SG}} = \Delta_{\text{Newt}}(1 + F(z)/2)$. If $|F(z)| < 0.05$, we simply use Newton's Method. In both cases, the fractional difference in step size is $\lesssim 10\%$. In the later case, we convert to Newton's Method for all future iterations without bothering to calculate $F(z)$ nor $p''(z)$.

To speed up the calculations even more, we stop calculating the stopping criterion once the Newton's method regime is reached. Usually this happens when we are already very

³In fact, we evaluate $|F(z)|^2 = (\text{Re}F)^2 + (\text{Im}F)^2 = \text{Re}(\overline{F}F)$, since this is substantially faster. The inequalities discussed in the text are then evaluated according to their squares.

close to the root, thus the value of stopping criterion will not change much. For the rare cases when this is not true, there is however a failsafe that returns the algorithm to the full Laguerre’s stage if the number of iterations in Newton’s stage exceeds 10.

Using this dynamic algorithm, instead of the full Laguerre method, leads to an increase in speed by about factor of 1.3. This algorithm is implemented in the `CMPLX_LAGUERRE2NEWTON` routine.

2.3. No Preference for Real Roots

There are many root-solver applications for which the distinction between real and complex roots is of fundamental importance. In particular, complex roots are often “unphysical”. Thus, generic root solvers may have a criterion that if the imaginary component of the root is within (a conservative version of) the numerical precision of zero, then it is simply set to zero. For example the *Numerical Recipes* `ZROOTS` sets this limit at about 10 times the naive numerical precision. Because the lens equation makes no fundamental distinction between real and complex roots, we do not include any such preference for real roots in the codes.

2.4. Stopping criterion

The problem of when to stop the Laguerre/Newton iterations is, in general, a subtle one. If the threshold is set conservatively above what is achievable given the numerical precision, then the result will not be as precise as possible. If set conservatively below, then the maximum precision will be achieved but at the cost that the condition is never met, so that iterations continue until some maximum allowed number is reached, which can burn a lot of computation time. This problem would seem to be solved by choosing the threshold at exactly what can be achieved. Unfortunately, this is not possible at the factor few level.

Adams (1967) derived a stopping criterion for polynomial root finding. It is based on a limit (E) on the value of polynomial that is indistinguishable from zero given the numerical error in polynomial evaluation. Whenever the value of the polynomial at a given point is below E ($|p(x)| < E$), one could stop looking for a better approximation of the root. A conservative formula for E gives this precision under the assumption that numerical errors add linearly in intermediate computation steps (Adams 1967, see Eq. (8)). However, from a statistical standpoint, these errors will tend grow in quadrature rather than linearly.

NR uses a simplified version of the Adams criterion, to enable faster calculations, by taking only the first component of the sum in Equation (10) in Adams (1967). We call this E_1 . We have found that this simplified criterion is a good approximation of the full

criterion, and that both will overestimate round-off errors in some fraction of cases. For the simplified criterion, the true underlying limit of precision will not be achieved in $\lesssim 10\%$ of cases. Typically, this limit lies a factor 8–100 below the criterion.

Figure 1 (upper panel) shows the value of the $|p(z)|/E_1$ at the step when the stopping criterion was satisfied. The Gaussian-like profile on the values of the $|p(z)|/E_1$ contains values corresponding to the real limit of accuracy. The tail going to higher values, up to the line $|p(z)| = E_1$, contains cases for which additional iterations would be warranted, since the true limit of precision was not achieved. The other panels in the Figure 1 show that one additional iteration is enough to get down to full numerical precision.

Therefore, in most algorithms in this paper, we employ the simplified Adams criterion with one additional step done after the criterion is satisfied. We skip this additional step in cases for which full precision is not essential to the result.

We make all our calculations in double precision. We understand “double precision” as being a 64 bit (8 byte) number described by IEEE 754 floating-point standard. This means 1 bit for sign, 11-bit exponent and 52-bit mantissa. Hence the fractional round-off error we should use in Equation (10) of Adams (1967) is $2 \cdot 10^{-15}$.

2.5. Subroutine CMLX_ROOTS_GEN

In CMLX_ROOTS_GEN we use the dynamic algorithm described in Section 2.2.3 for the “robust” part of calculations. Because this algorithm incorporates Newton’s method, in a small number of cases (e.g., $\sim 10^{-5}$ in our experimentation on 5th order polynomial) it will fail to converge in the prescribed number of iterations. Our algorithm therefore checks for these failures, and then employs a full Laguerre search starting from point (0,0), which finds correct root more robustly (at the cost of greater computation time).

We note that in the last iteration, there is no need to divide the 1st order polynomial by the last root that was found. Additionally, the last root can be easily found by using Viète’s formula (Viète 1579) for quadratic equation. These two optimization, skipping one root solver and two divisions, leads to 10% gain in the performance when the algorithm is used to solve a 5th order polynomial.

During “polish”, when we try to find a more accurate position of the root based on the full (not divided) polynomial, we can use Laguerre’s method or Newton’s method. Since usually polishing takes only one step, this choice does not have a large impact on the overall speed of the algorithm.

There is a flag in the argument list of the `CMPLX_ROOTS_GEN` routine called “`use_roots_as_starting_points`”. If the routine is started without any knowledge of the root positions, this flag can be set to “`false`”, so all components of the “roots” array will be reset to $(0, 0)$ inside the routine. On the other hand, if one knows the approximate position of at least one root, the value(s) should be put at the end of the “roots” array and the flag should be set to “`true`”. (Unknown values should be reset to $(0, 0)$.) This will help the routine to find all roots faster. It is important to initialize the whole “roots” array prior calling the routine unless the “`use_roots_as_starting_points`” flag is set to “`false`”.

This routine, with the use of the mentioned flag, can be used to robustly find all roots of a polynomial that was created from an already solved polynomial by making small changes in its coefficients. This is faster than starting all searches from zero. Robustness comes of course with a price of speed, when compared to pure root polishing. In Section 3 we present a method that takes advantage of pure root polishing speeds, but employs a series of checks to ensure robustness in the particular problem of the binary lensing.

Subroutine `CMPLX_ROOTS_GEN`, provided in the source code, implements the robust, general complex polynomial solver with the use of all optimizations mentioned above. It is consistently faster than the NR implementation, `ZROOTS`. For 5th order polynomials we see a speed increase⁴ by a factor of 1.6–1.7.

3. Further optimization for binary lensing

There are three broad regimes in which the binary-lens solver must be applied. The first is the point-source regime, for which the magnification can be approximated as constant over the source. This requires only a single call to the lens solver. In the second regime, the source suffers moderate differential magnification due to the proximity of a caustic. This is generally handled by the hexadecapole approximation (Pejcha & Heyrovský 2009; Gould 2008), which requires 13 calls to the lens solver. Finally, if the source actually straddles a caustic, or if it is sufficiently close that differential magnification is severe, then much more computationally intense methods are required. These generally fall into two classes: inverse ray-shooting (Kayser et al. 1986; Wambsganss 1990, 1999) and contour integration (Gould

⁴We perform all speed tests using the Intel Fortran compiler (ifort) with “`-O2`” optimization flag. We noticed that the algorithms discussed in this paper work two times faster when compiled with ifort rather than Gnu Fortran Compilers (gfortran or g77). Thus, we encourage users to test speeds of their programs against different compilers and compilation flags. This exercise can lead to comparable gains in performance relative to the matters discussed in this paper.

& Gauchere 1997, e.g.). All methods require calls to the lens-solver, but in the ray-shooting approach the majority of the computation, as well as the precision of the result, depends primarily on another numerical operation, namely evaluating the lens equation. Therefore, it is for contour integration, which uses Stokes’ Theorem, and the hexadecapole method, that calls to the lens solver are the rate-limiting step. And it is for these approaches that numerical errors in the lens solver propagate into the final result. Since these are widely used, and since such computations are often quite lengthy and usually time-critical, it is worth an effort to optimize the lens solver.

An efficient algorithm tailored specifically for root polishing is crucial for contour integration as well as for the hexadecapole method, where one needs to locate roots multiple times at close-by positions of the source.

3.1. Simplifying Features of the Binary Lens Equation

The binary lens equation has two simplifying features relative to general complex polynomials. First, it is fifth-order, and therefore just “one step” away from being susceptible to analytic solution (Abel 1824). Second, two of the five roots are always isolated. That is, when the source approaches a caustic (or cusp), then two (or three) of the roots can be very close together, but the remaining roots are always well separated from these and from one another. Together, these two features enable a very different approach.

3.2. New Algorithm: Concise Overview

Our original idea was to improve the root polishing algorithm by identifying the two isolated roots and solving these by Laguerre’s Method. Then we would divide these out and solve for the remaining three roots using the cubic equation. Because the first two roots are isolated, they are very precisely determined. Therefore the normal concerns about introducing numerical noise by dividing out roots, which do very much apply to closely-spaced roots, are greatly mitigated.

However, first we found that Newton’s method is faster than Laguerre’s, and is also quite robust for these two isolated roots when they are already approximately known. Second, somewhat surprisingly, we found that the third most isolated root was more precisely located by Newton’s Method than by the cubic equation (though not by much). Finally, we found that the remaining two (closest) roots could be found roughly 3 times faster than Newton’s Method by dividing out the first three roots and solving the resulting quadratic equation, thus, effectively increasing the speed of polishing by 25%. The quadratic equation is also slightly more stable than either Newton’s Method or Laguerre’s method for these two roots

in cases when a third root is nearby.

Of course, to apply this approach, it is necessary to know in advance which are the isolated roots. In some cases these are indeed known because one has just finished solving the lens equation at a very nearby source position. But for cases in which it is not known, we simply apply the iterative root searching algorithm `CMPLX_LAGUERRE2NEWTON` (see Section 2.2.3) to find two (now arbitrary) roots, and use the cubic equation to find the rest. The two closest roots can then be easily identified.

4. Detailed Description

The algorithm works in 2 modes: “robust”, which can be started without any knowledge of the roots, and “polish”, which is meant to be used with approximately known roots, e.g., with roots that come straight from “robust”, or come from the previous solution to a similar polynomial. The flag “polish_only” controls the behavior of the subroutine. We describe both modes in turn, and discuss the failsafe measures that are incorporated in “polish”.

We still must start with Laguerre’s Method in “robust” mode (Section 4.1) because it is guaranteed to find a root whereas Newton’s Method is not. But because of the simplifications of the Laguerre calculation, as described in Sections 2.1 and 2.4, as well as use of the new intermediate method described in Section 2.2, Equation (8), we are able to cut the time spent in “robust” by a factor of 1.6-1.7 in most cases.

4.1. Robust

This mode assumes that absolutely nothing is known of the location of the roots. We begin by searching for the first root by applying the previously described general method seeded at (0,0). We then divide the original polynomial by the root and repeat the process to find a second root. (Note that although this process is called root “division”, it is actually composed of only multiplications and additions, and is therefore very fast compared to Laguerre.) If one or both of these roots is close to other roots, then root division may introduce substantial errors because these root positions can be determined only to a precision equal to the square-root or cube-root of the underlying numerical precision. See Section 5. So for example, for “complex*16” precision (with 52 mantissa bits), the location of three close roots can have errors up to $\sim 2^{-52/3} \sim 10^{-5}$. Nevertheless, such errors are quite unimportant at this stage because the isolated roots will still be approximately located at the following step when the remaining cubic equation is solved analytically. This method is guaranteed to find five different roots, unless the source happens to land “exactly” (to numerical precision) on a caustic. Note that the only real information that will be extracted from this aspect of the subroutine is the approximate position of the three most isolated roots. The five roots are

then fed to the “polishing” routine.

As will become clear, it is essential that the fourth and fifth root in this list be the two closest roots. So we must, at a minimum, locate this pair. An efficient algorithm for doing so is located in subroutine “find_2_closest_from_5”. However, in fact we strictly order *all* roots according to which is most isolated. The distances between all pairs of roots are calculated. The one whose minimum distance to other roots is the largest is declared most isolated. If two roots are tied (as will often be the case), the honor goes to the root whose second-closest neighbor is more distant. Then other roots are ranked in a similar fashion. (See subroutine “sort_5points_by_separation”). This is useful for some applications, but increases the run-time of this sub-algorithm by about 5%. Users may therefore decide to substitute the simpler algorithm for finding the closest root pair.

4.2. Polish

If “polish_only” is set to true, then the algorithm described below will be acted upon immediately, starting from the root positions provided in the input array. In the opposite case, first the subroutine will carry out the steps from “robust” (Section 4.1) and subsequently send the resulting list of roots to be polished.

The polish algorithm initially assumes that the first three roots in the input list are the most isolated (but see immediately below). Then it polishes these using Newton’s Method. It then successively divides the polynomial by these three polished roots, and solves the remaining quadratic equation to find the last two.

Then it checks to make sure that these are in fact the closest pair out of all roots. Of course, if the trial roots come from the “robust” algorithm, then this will almost certainly be the case. And if the trial roots came from a neighboring position on the contour, then it will still almost always be the case. Nevertheless, as the source position is moved around a contour, the pairs of roots that had been closest move apart, while others move close together. Therefore, depending on exactly how the contour is sampled, such switches in closest pair may occur a few percent of the time.

When this does occur, the roots are completely reordered according to degree of isolation (as at the end of Section 4.1) and are re-polished. A flag “first_3_roots_order_changed” is set to notify the calling routine that there has been a reordering. This is because, for contour integration, the calling routine must match roots from one step to the next. If there has been no reordering, then the first three roots can be assumed to be automatically matched, and it is only necessary to check the final two. But if the roots have been reordered, then full matching of all roots is required.

4.3. Failsafe for Polishing

As described above, if the algorithm is started in “robust” mode, the roots will be found with maximum precision, beginning with no information. If the flag is set to do only “polish”, then it will achieve the same result with less computation, but only on the condition that the input roots are roughly correct. If incorrect roots are fed into the “polish” routine, it can in principle fail catastrophically in one of two following senses. First, it may fail to find a root after the prescribed maximum number of iterations (50) because it is using Newton’s Method rather than Laguerre’s, which is more robust. Second, two of the first three (“isolated”) root searches may yield the same root. Therefore, if “polish” finds that the last two roots are not the closest, it simply restarts algorithm in the “robust” mode.

Only when the two last (“closest”) roots after polishing stay the closest, does the algorithm return without any additional action. If it was called with the “polish_only” flag, it also informs the calling routine that there was no reordering of the first 3 roots.

Sending back to “robust” may seem very expensive, but this is the only way we can ensure that 5 distinct roots will be returned by the algorithm. Polishing always implies some level of risk of collapsing two roots to the same point, in which case the resulting pair of roots will be the closest pair and the failsafe will be triggered. In the rare case that subsequent polishing, after “robust”, is not successful, the algorithm decides to return unpolished roots, since clearly there are some numerical problems, and better accuracy might be hard to achieve.

We note that after the failsafe procedure is triggered, we can still use some of the information gathered during the previous stage to facilitate calculation in robust. We simply copy two roots found in the polishing stage that did not end up in the closest pair, and use those as a starting points of the two searches in “robust”. This makes the failsafe algorithm much less time demanding than starting the whole routine all over again.

This failsafe feature, which has zero cost for normal polishing, means that the “polish” mode can be used for example for point-lens portions of the light curve. Typically, these contain a series of points that are close enough together that only polishing is required. But for the cases that this proves not to be so (for example, the source has passed over a caustic between one night and the next), the algorithm itself can figure this out and correct it. The cost is an extended failed polishing, which can take longer than a simple call to “robust”. However, if it is known that this involves only a few percent of points, this cost would be substantially smaller than the savings.

As we show in Section 5, for very specific applications, it is warranted to do all polishing numerically rather than using the quadratic equation. Comments in the code can point the

reader to specific changes for a “polish only with Newton” scenario. Notably, there is one more check in the failsafe procedure required in this case. After polishing, one should check whether the distance between closest pair of roots is smaller than some threshold (e.g. 10^{-9}). If so, this could suggest that two of the roots collapsed during polishing, and polynomial should sent to “robust”.

4.4. Application to Hexadecapole

As mentioned in Section 3, the hexadecapole approximation is used when the source is too close to a caustic to apply the point-source approximation, but not so close as to require a full finite-source calculation. It requires 13 calls to the lens solver with a specified geometric distribution (Gould 2008). All but the first of these can safely use “polish”, seeding it with the first set of roots found for the source center. The combination of substituting “polish” for “robust”, combined with the improved speed of “polish” described in Section 4.2, implies an overall factor of 2-3 in the root finding speed for hexadecapole calculations. Of course, these calculations still require calculation of the magnifications for each of the 13 calls; these are however much cheaper than locating the roots.

5. Analytic Error Estimates

Here we derive analytic estimates of the precision achievable when two or three roots are close together and test these numerically. Of course, since the complex plane has no intrinsic scale, one must define “close” relative to something. In this case, we mean close relative to some additional (third or fourth) root.

5.1. Two Close Roots

Suppose that a cubic has roots $\pm c$ and d , where $|d| \sim 1$ and $|c| \ll 1$. The polynomial will then have the form

$$f(z) = (z^2 - c^2)(z - d) = z^3 - dz^2 - c^2z + c^2d \quad (14)$$

We then evaluate the isolated root d using Newton’s (or Laguerre’s) Method, but inevitably make an error ϵ which is of order of the numerical precision. Next we divide $[z - (d + \epsilon)]$ into $f(z)$ so as to obtain the quadratic equation

$$\frac{f(z)}{z - (d + \epsilon)} = z^2 + \epsilon z - (c^2 - \epsilon d) \quad (15)$$

and a possible remainder, which we would be forced to ignore. This can be solved using the quadratic equation

$$z = -\frac{\epsilon}{2} \pm \sqrt{c^2 - \epsilon d} \quad (16)$$

where we have already made use of the fact that $|\epsilon| \ll |d|$ to drop a term of order ϵ^2 . Considering the two relevant regimes, ($|c^2| \gg |\epsilon d|$) and ($|c^2| \ll |\epsilon d|$) and now explicitly evaluating $|d| = 1$, we see that the error in c is of order

$$|\delta c| \sim \frac{|\epsilon|}{2|c|} \quad (|c| \gg \sqrt{|\epsilon|}); \quad |\delta c| \sim \sqrt{|\epsilon|} \quad (|c| \ll \sqrt{|\epsilon|}). \quad (17)$$

Note, however, that while the errors in the values of these roots scale $\sim |\epsilon|^{1/2}$, the error in the midpoint between the two roots scales $\sim |\epsilon|$, i.e., it is much smaller.

5.2. Three Close Roots

Suppose that a quartic has roots $ce^{im2\pi/3}$ and d , where again $|d| \sim 1$ and $|c| \ll 1$ and where $m = 0, \pm 1$. The polynomial will then have the form

$$f(z) = z^4 - dz^3 - c^3z + c^3d. \quad (18)$$

We again divide by the isolated root, which again is in error by ϵ , and obtain the cubic

$$\frac{f(z)}{z - (d + \epsilon)} = z^3 + \epsilon z^2 + \epsilon dz - (c^3 - \epsilon d^2). \quad (19)$$

This yields roots

$$z = -\frac{\epsilon}{3} + e^{im2\pi/3}(c^3 - \epsilon d^2)^{1/3}, \quad (20)$$

which has analogous error properties to Equation (17),

$$|\delta c| \sim \frac{|\epsilon|}{3|c|^2} \quad (|c| \gg |\epsilon|^{1/3}); \quad |\delta c| \sim |\epsilon|^{1/3} \quad (|c| \ll |\epsilon|^{1/3}). \quad (21)$$

Thus, for “complex*16” precision, which contains 52 bits of mantissa, the limiting precision for very close roots is $10^{-7.7}$ for two close roots (near a caustic) when calculated using the quadratic equation, and $10^{-5.2}$ for three very close roots (near a cusp), when obtained by solving the cubic equation.

5.3. Numerical tests

We test the accuracy in calculation of close-by roots numerically. Figure 2 shows typical errors in position estimation of the two close roots when solving the quadratic equation, as well as errors in positions of three close-by roots when solving the cubic equation, as a function of distance between the roots (c). These behave as described by Equations (17) and (21), with two regimes clearly visible. For a characteristic scale of the problem $d \sim 1$, the transition occurs at $\epsilon^{1/2}$ and $\epsilon^{1/3}$, respectively.

We also find (not shown) that pushing away one of the three close-by roots results in monotonically increasing accuracy, which bridges the relations for three and two close-by roots.

The choice of origin of the coordinate system has a very low impact on the accuracy achievable by the quadratic and cubic solvers. This is because the biggest error comes from the division of the polynomial by imperfectly known roots. However, we noticed that in the case of iterative numerical methods, like Newton’s Method, used on the undivided polynomial, the accuracy depends on the distance of the close-by roots from the origin of coordinates (o), scaling as $\epsilon \rightarrow \epsilon \cdot o^2$ in Equation (17), for the case of 2 nearby roots, and as $\epsilon \rightarrow \epsilon \cdot o^3$ in Equation (21), for the case of 3 nearby roots. Figure 3 shows how the accuracy improves when we bring the two close roots nearer to the origin of the system and solve them with Newton’s Method (with undivided 5th order polynomial) rather than the quadratic equation, obtained by dividing the polynomial by 3 known roots. Similarly, Figure 4 shows the case of 3 close-by roots.

Note that to take advantage of this feature, one should think in advance about where the close-by roots are expected, and then shift the coordinate system to this point. However this may not always be possible. Also, this level of accuracy is not usually required for the close-by roots in microlensing. This impacts only a small number of source positions that are very close to the caustic. The gain in speed achieved by the use of the cubic and quadratic solvers for all calculated positions, in most cases, is more important than the slight loss of the accuracy for this small subset of points. However, for certain applications, taking advantage of this behavior can be fruitful.

We checked that we can recover the accuracy of numerical estimation of the roots on the undivided polynomial, when we use quadruple precision for polynomial division. Before division, we perform one additional Newton step in quadruple precision on each root that will be used in the division. The resulting quadratic equation can be still solved with double precision without loss of accuracy.

We thank Google Books for digitizing the World’s knowledge and making it public for free, making possible scientific works, such as this one. We thank Wikipedia for providing quick references to historical literature and concise explanations of many mathematical methods. In particular, Equation (1), which played a major role in the genesis of this work, was derived from a closely related equation found in Wikipedia. This work was supported by NSF grant AST 1103471.

A. Origins of “Newton’s Method”

We began our investigation of the historical root-solving literature simply to identify proper references to Newton, Halley, and Laguerre. This proved to be a non-trivial task despite the considerable help provided by several Wikipedia entries and by Google Books. In the course of reading the various sources, simply to verify that we had found the right ones, we discovered that the method laid out in a work Newton claims⁵ to have written in 1667-1671 (Newton 1671), does not contain “Newton’s Method”, nor anything approximating it. Moreover, in this work, Newton exhibits the deepest confusion over the relationship between the method he presents and the mathematics that are actually required for “Newton’s Method”.

Since this method is sometimes referred to as “Newton-Raphson”, we initially thought that these shortcomings would be corrected by Raphson (1690). However, we found that while Raphson’s work (which appears to be completely independent of Newton) is indeed superior to Newton’s, it still does not approximate “Newton’s Method” anywhere closely enough to designate Raphson as its discoverer.

Eventually, we came upon a biography of Thomas Simpson, which stated that it was Simpson who discovered “Newton’s Method”, while Newton and Raphson had found only algebraic prescriptions for polynomial root solving that did not involve calculus. This point is indeed central, although the gulf between Newton (and/or Raphson) and “Newton’s Method” actually runs much deeper. The biography gave “1740” as the date of Simpson’s discovery, but did not cite a work. In fact, Simpson wrote two works in that year, each several hundred pages, and it proved fairly difficult to locate his description of “Newton’s Method” because it is given without benefit of equations (displayed or otherwise). Nevertheless, Simpson’s description is quite clear (and quite short). Moreover, Simpson clearly and correctly states that this is a “new method” (Simpson 1740).

Armed with “keywords” from Simpson’s work, we were able to locate an article by Kollerstrom (1992) in the British Journal for History of Science, which makes the case for Simpson’s discovery quite clearly:

What is today known as “Newton’s method of approximation” has two vital characteristics : it is iterative, and it employs a differential expression. The latter is simply the derivative $f'(x)$ of the function, resembling a Newtonian fluxion in being based upon a theory of limits but not conceptually identical with it. The method uses the fundamental equation $[\Delta\alpha = -f(\alpha)/f'(\alpha)]$ repetitively, inserting at each stage the (hopefully) more accurate

⁵There is no basis for doubting this claim, but neither is there any independent evidence supporting it.

solution. This paper will argue that neither of these characteristics applies to the method of approximate solution developed by Newton in De analysi, which also appeared in his De methodis fluxionum et serierum infinitarum, and that the method of approximation published by Joseph Raphson in his Analysis aequationum universalis of 1690 was iterative – indeed was the first such method to be iterative – but was not expressed in derivative or fluxional terms.

We fully agree with this assessment, and we strongly urge the interested reader to study the entire article, which also contains an excellent discussion of how the myth of Newton’s discovery of this method was born and how it “blossomed” over several hundred years.

However, we would also like to make two points that go beyond Kollerstrom’s analysis. First, in addition to the two distinguishing characteristics of “Newton’s Method” identified by Kollerstrom (“iterative”, and “differential calculus”), there is a third characteristic that we also regard as essential: the method is applicable to an extremely broad class of functions, essentially all differentiable functions. Now, it is well-known that Newton (and Raphson) only applied their methods to polynomials. Moreover, it can hardly be regarded as a shortcoming that they did so, since the general theory of functions was so weakly developed at this point. But what does not seem to have been previously appreciated is that their *methods* are so specifically designed to solve polynomials that they make use of features of polynomials that do not hold for almost any other function. Therefore, these methods cannot be generalized to all functions in anything like their present forms. That is, they cannot be generalized without specifically reformulating them on the basis of differential calculus.

This brings us to the second point. Newton’s exposition of his method makes explicit (in so far as lacunæ can be explicit) that he does not understand the relationship between the method he presents and differential calculus. This is particularly striking to the modern reader, who is aware that Newton invented calculus, that “Newton’s Method” is about the simplest application of differential calculus imaginable, and that Newton was one of the most brilliant mathematicians and physicists of all time, and so naturally assumes that Newton *must* have been cognizant of this relation, whether explicitly or implicitly. But this is clearly not the case. In his *Method of Fluxions*, Newton presents his extremely clunky, algebraic method for solving polynomials, and then immediately after finishing this section, begins a new section on derivatives, which is quite elegant, including a diagram of a function, with its derivative represented as a tangent, just as appears in countless modern textbooks. Indeed, this figure could have just as easily been used to illustrate “Newton’s Method” as it is presently understood. The fact that it was not, virtually proves that Newton did not understand this method in anything like its modern form.

B. Notes on copyright of Numerical Recipes

Our original motivation to write our own root-solving algorithm was not to improve upon ZROOTS, which like Bozza (2010) we believed to be essentially impossible. Instead, we merely sought to create a root solver that was sufficiently different from ZROOTS that we could make publicly available some binary-lens code without fear of lawsuit by the authors of *Numerical Recipes* (NR) on charges of copyright violation.

The actual outcome of this effort, i.e., new algorithm with dramatically better performance, confirms (if such confirmation were needed) the value of copyright and patent laws in stimulating intellectual innovation.

Nevertheless, as important as this stimulus is, we believe that copyright protection of NR algorithms is at this point, on balance, a substantial obstacle to scientific progress. Teuben et al. (2012) have recently argued that “code discoverability” is becoming increasingly important to the fundamental criterion for acceptance of scientific conclusions: “reproducibility of results”. This concern can be interpreted narrowly in terms of allowing independent groups to run the same, or slightly modified code, but also more broadly, as making available the entire intellectual basis of a numerical discovery in order to enable and stimulate further advances on the topic. We urge the interested reader to review this article, which concisely addresses many of the arguments (aka excuses) given by numerical researchers for not publishing their codes.

However, at least in astronomy, many codes contain commercial subroutines, and the majority of these are probably from NR. These cannot at present be made public without violating copyright laws.

The general problem here is quite complex since it arises from a conflict of cultures between commercial and non-profit activities in our society. Code development, particularly of the high quality represented by NR, is not free. It must be supported either by universities and public research organizations, or by private ventures that anticipate revenues to cover the labor invested in writing and distributing the code, as well as normal profit on invested capital. The former route may be regarded as a “natural outcome” of research activity, but in fact is heavily influenced by the whims of faculty-search and P&T committees. NR is an example of the latter route. One may argue about whether the degree of remuneration has been appropriate to the effort, but certainly the results of these efforts have had far more positive impact than many activities that are more highly rewarded (Morgenson 2011)

A general solution to this problem is beyond the scope of the present work. One solution might be to include algorithm purchases on grant proposals (similar to how computer purchases are presently handled), with the stipulation that both the purchased algorithms

and the new algorithms covered under the award be made public for applications restricted to those directly traceable to the funded work. The mere statement of this proposal conveys how difficult it would be to properly formulate in practice. This is why we do not attempt a general solution here.

However, most immediately, the problem could be solved if the NR authors agreed to make their algorithms available at no cost to non-profit users, with the stipulation that this did not extend to the commercial sector. It might be possible to do this through specifically worded licensing agreements. We urge the NR authors to consider such an approach.

C. List of subroutines in the open-source code

We provide two versions of the codes, written in Fortran 90 and Fortran 77, for convenience of the users. Both versions have the same set of subroutines, and their output is the same. Sources are located in files: `cmplx_roots_sg.f90` and `cmplx_roots_sg_77.f`. In the package we also provide files called `LICENSE` and `NOTICE` containing open-source licensing details. Table 1 list all subroutines provided in the codes with short explanations of their use. We also list all arguments required to call the routines, as well as their role, type and if they are meant for input only (“in”), output only (“out”) or both (“in/out”).

Table 1:: List of subroutines.

- 1 **CMPLX_ROOTS_GEN** – General complex polynomial solver
 - `roots` – array of roots, `length=degree` – (`complex*16`, in/out)
 - `poly` – input polynomial, array `length=degree+1`, `poly(1)` is a constant term – (`complex*16`, in)
 - `degree` – degree of input polynomial and size of `roots` array – (`integer`, in)
 - `polish_roots_after` – turns on polishing, uses undivided polynomial after all roots are found – (`logical`, in)
 - `use_roots_as_starting_points` – if set to `.false.` then `roots` array will be reset to (0,0). Otherwise the values in `roots` will be used as starting points for each root – (`logical`, in)

- 2 **CMPLX_ROOTS_5** – 5th order polynomial solver for binary lens equation
roots – array of roots, length 5 – (**complex*16**, in/out)
first_3_roots_order_changed – output flag showing if reordering of roots occurred – (**logical**, out)
poly – input 5th order polynomial, array length 6 – (**complex*16**, in)
polish_only – if **.true.** then “robust” is skipped and algorithm goes to “polish” starting from values given in **roots** array – (**logical**, in)
- 3 **SORT_5_POINTS_BY_SEPARATION** – This sorts array of 5 points to have the 1st point most isolated, and 4th and 5th being closest
points – array to sort, in place – (**complex*16**, in/out)
- 4 **SORT_5_POINTS_BY_SEPARATION_I** – This sorts array of 5 points by separation, but returns only indices of a sorted array
sorted_points – array of 5 indices that would sort array, 1st index shows the position of most isolated point from array **points** – (**integer**, out)
points – array of points to be sorted, length=5 – (**complex*16**, in)
- 5 **FIND_2_CLOSEST_FROM_5** – Routine to find closest pair from set of 5 points.
i1 – index of the one component of the closest pair from array **points** – (**integer**, out)
i2 – index of the second component of the closest pair from array **points** – (**integer**, out)
d2min – square of the distance between closest pair of points – (**real*8**, out)
points – array of points, length=5 – (**complex*16**, in)
- 6 **CMPLX_LAGUERRE** – Single root finding routine which uses Laguerre’s Method.
poly – input polynomial, array of length $\geq \text{degree}+1$ – (**complex*16**, in)
degree – degree of polynomial – (**integer**, in)
root – on input: a starting point for the algorithm, on output: a found root – (**complex*16**, in/out)
iter – number of Laguerre iterations done before returning – (**integer**, out)
success – success flag, if number of iteration is higher than specified number **.false.** is returned – (**logical**, out)

- 7 **CMPLX_NEWTON_SPEC** – Single root finding routine which uses Newton’s Method. With modification that stopping criterion is calculated only every 10th step.
- poly** – input polynomial, array of length $\geq \text{degree}+1$ – (complex*16, in)
degree – degree of polynomial – (integer, in)
root – on input: a starting point for the algorithm, on output: a found root – (complex*16, in/out)
iter – number of Newtons iterations done before returning – (integer, out)
success – success flag, if number of iteration is higher than specified number .false. is returned – (logical, out)
- 8 **CMPLX_LAGUERRE2NEWTON** – Single root finding routine which uses new dynamic method with three regimes: Laguerre, SG and Newton.
- poly** – input polynomial, array of length $\geq \text{degree}+1$ – (complex*16, in)
degree – degree of polynomial – (integer, in)
root – on input: a starting point for the algorithm, on output: a found root – (complex*16, in/out)
iter – number of total iterations done before returning – (integer, out)
success – success flag, if number of iteration is higher than specified number .false. is returned – (logical, out)
starting_mode – if 2 – starts with Laguerre’s method, if 1 – starts with SG, if 0 – starts with Newton’s method – (integer, in).
- 9 **SOLVE_QUADRATIC_EQ** – The analytic solver of quadratic equation.
- x0** – first root – (complex*16, out)
x1 – second root – (complex*16, out)
poly – input polynomial of degree=2, length ≥ 3 , $\text{poly}(1) + \text{poly}(2) x + \text{poly}(3) x^2$ – (complex*16, in)
- 10 **SOLVE_CUBIC_EQ** – The analytic solver of cubic equation using Lagrange’s Method.
- x0** – first root – (complex*16, out)
x1 – second root – (complex*16, out)
x2 – third root – (complex*16, out)
poly – input polynomial of degree=3, length ≥ 4 – (complex*16, in)

- 11 `DIVIDE_POLY_1` – Division of the polynomial by monomial (x-p).
`poly_out` – resulting polynomial after division with `degree=degree-1`, array
of length $\geq \text{degree}-1$ – (`complex*16`, `out`)
`remainder` – remainder from the division – (`complex*16`, `out`)
`p` – coefficient in (x-p), i.e., in monomial by which `poly_in` is divided –
(`complex*16`, `in`)
`poly_in` – input polynomial of a `degree=degree`, array of length $\geq \text{degree}$ –
(`complex*16`, `in`)
`degree` – degree of the input polynomial – (`integer`, `in`)

REFERENCES

- Abel, N., H. 1824, Mémoire sur les équations algébriques, où on démontre l'impossibilité de la résolution de l'équation générale du cinquième degré, Grøndahl, Christiania, Oslo
- Abel, N. H. 1826, "Beweis der Unmöglichkeit, algebraische Gleichungen von höheren Graden als dem vierten allgemein aufzulösen." J. reine angew. Math. 1, 65. Reprinted in Abel, N. H., 1881, Oeuvres completes de Niels Henrik Abel, ed. L. Sylow and S. Lie, Christiania Imprimerie De Grøndahl & Son, Oslo.
- Adams, D. A., 1967, Communications of the ACM, Volume 10, Issue 10, Oct. 1967, p. 655
- Bozza, V. 2010, MNRAS, 408, 2188
- Dominik M., 1993, Effiziente Methoden zur Invertierung der Gravitationslinsengleichung und zur Analyse von Bildern ausgedehnter Quellen, Diploma thesis, Universität Dortmund
- Dominik, M. 1998, A&A, 333, L79
- Gould, A., & Gaucherel, C. 1997, ApJ, 477, 580
- Gould, A. 2008, ApJ, 681, 1593
- Halley, E., 1694, A new and general method of finding the roots of equations, Phil. Trans. Roy. Soc. London, vol. 18, p. 136

- Kayser, R., Refsdal, S., & Stabell, R. 1986, A&A, 166, 36
- Kollerstrom, N., 1992,
- Laguerre, E., N., 1880, *Nouvelles Annales de Mathématiques*, 2e série, t.XIX, 87-103, reprinted in Laguerre, E., N., 1898, *Oeuvres de Laguerre*, eds. MM. CH. Hemrite, H. Poincaré and E. Rouché, Académie des sciences (France), pp. 87-103
- Morgenson, Gretchen, 2011, *Reckless Endangerment: How Outsized Ambition, Greed, and Corruption Led to Economic Armageddon*, New York, Times Books
- Newton, I., 1671, *Methodus fluxionum et serierum infinitarum cum ejusdem applicatione ad curvarum geometriam*, published in English in Newton, I., 1736, *The methods of fluxions and infinite series with its application to the geometry of curve-lines*, ed. and transl. by John Colson, London, Henry Woodfall, and in Latin: Newton, I., 1744, *Methodus fluxionum et serierum infinitarum*, *Opuscula Mathematica Philosophica et Philologica*, ed. Johann von Castillon, Apud Marcum-Michaelem Bousquet & Socios., Volume 1, Op. II, pp. 29-200
- Pejcha, O., & Heyrovský, D. 2009, ApJ, 690, 1772
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1992, *Numerical Recipes*, Cambridge: Cambridge Univ. Press
- Raphson, J., 1690, *Analysis Aequationum universalis*, London
- Simpson, T., 1740, *Essays on several curious and useful subjects: in speculative and mix'd mathematicks*, printed by H. Woodfall, jun. for J. Nourse, Section 6, pp. 81-86
- Teuben, P., Allen, A., Nemiroff, R. J., & Shamir, L. 2012, arXiv:1202.1026
- Viète, F., 1579, *Opera mathematica*. Reprinted Leiden, Netherlands, 1646
- Wambsganss, J. 1990, Ph.D. Thesis, Ludwig-Maximilians-Univ., Munich, Fakultät für Physik.
- Wambsganss, J. 1999, *Journal of Computational and Applied Mathematics*, 109, 353
- Witt, H. J., & Mao, S. 1995, ApJ, 447, L105

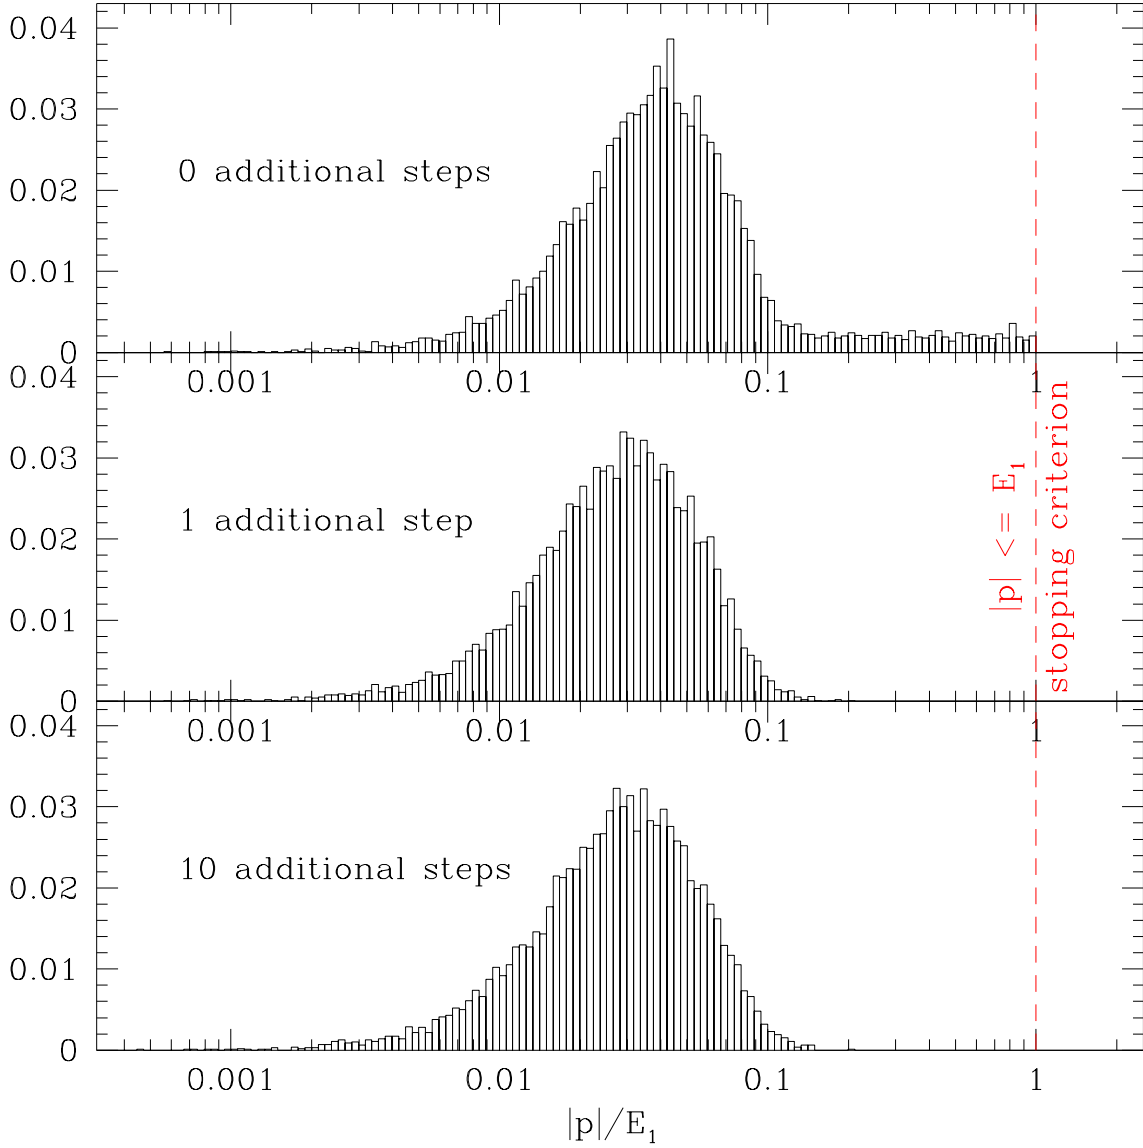


Fig. 1.— Simplified Adams (1967) criterion (E_1) versus achievable numerical precision. The three panels contain histograms of polynomial values evaluated at the calculated root, normalized to the scale of the stopping criterion ($|p|/E_1$), at three moments in the course of iteration. The first panel shows that $\lesssim 10\%$ of points do not reach final precision at the step when stopping criterion is met (and zero additional iterations are performed). The middle panel shows that one needs one additional iteration past this to achieve full possible precision. The third panel, which contains values of polynomial evaluated at the root approximation found after 10 additional iteration past the time stopping criterion was met, is given for comparison.

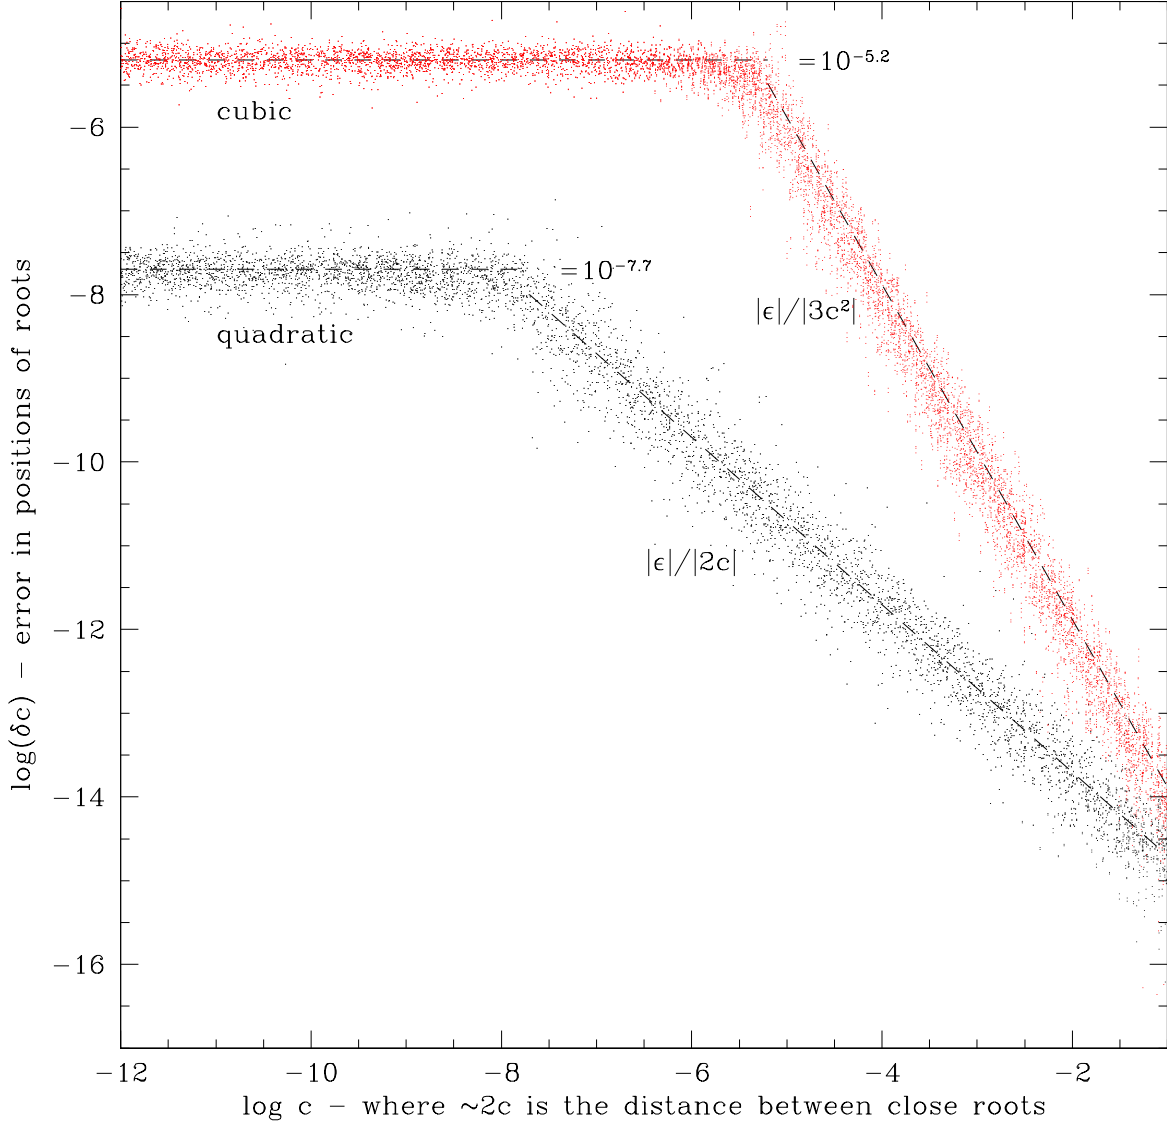


Fig. 2.— Result of numerical simulations of errors made when two close roots of a polynomial are solved using the quadratic equation (black points), and three close roots are solved using the cubic equation (red points). All other roots are at a distance of order unity from the origin and are found using numerical methods, then divided out leaving a quadratic or cubic equation. The dashed lines show the analytic estimates presented in Equations (17) and (21), respectively. The characteristic distance between close-by roots is $2c$; δc is the measure of error one would make in the position of the root, when solving the quadratic or cubic equation, which are results of division of the original polynomial by other roots. Newton’s Method used on the undivided polynomial leads to the same errors if the close-by pair or triple of roots is located at distance ~ 1 from the origin of the coordinate system.

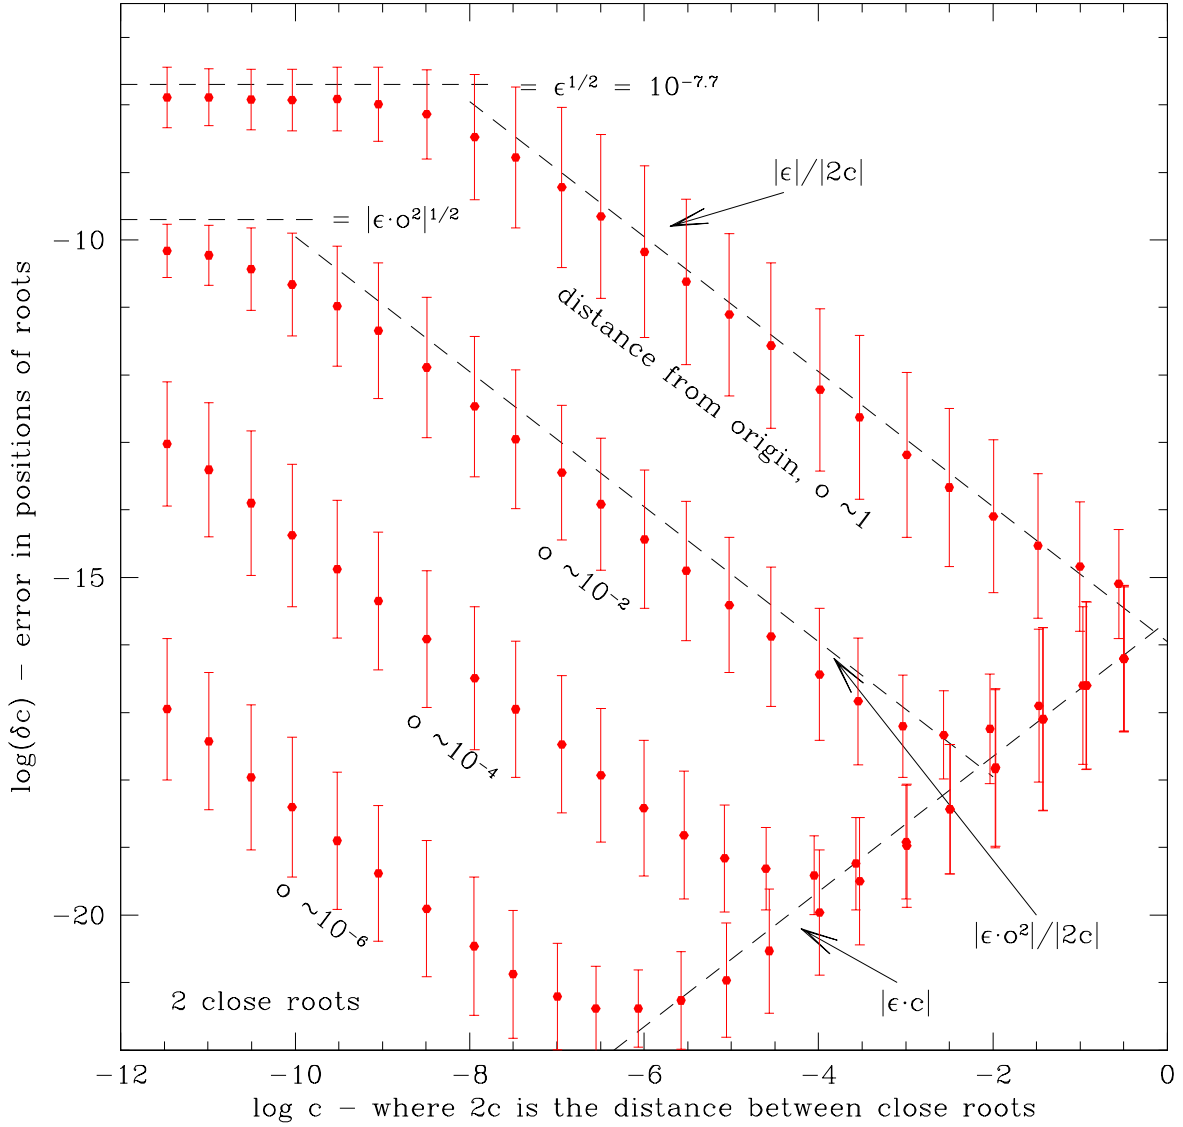


Fig. 3.— Illustration how errors of 2 close-by roots scale with the distance between those roots. Upper dashed lines show analytic estimates of the errors one should expect when solving the two last roots of a polynomial with a quadratic solver - Equation (17), see also Figure 2. Points with errorbars show the result from simulations when the two last roots were found numerically using the undivided polynomial, for cases when the close pair of roots was located at distances 1, 10^{-2} , 10^{-4} , and 10^{-6} from the origin of coordinate system. We see that when the characteristic distance from the origin of the last pair of close-by roots is close to unity, $o \sim 1$, the accuracy of the numerical algorithm is well approximated by the analytic result. When the distance o is smaller, the accuracy of the numerical results scales with o^2 in the case, when 2 roots are close. ($|\epsilon \cdot c|$ line marks the limit of double precision representation of a number).

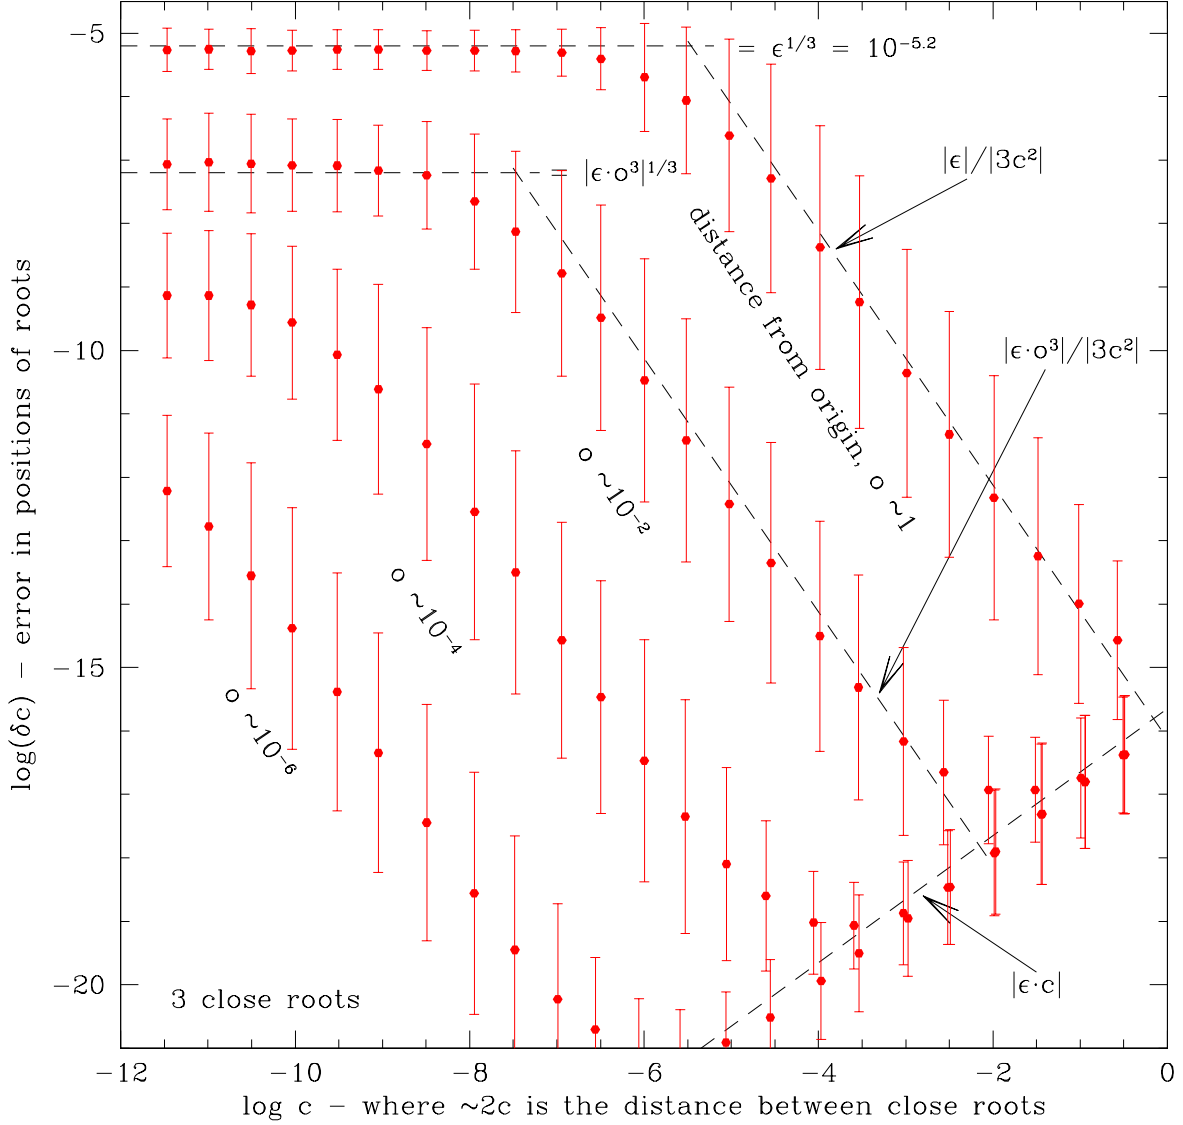


Fig. 4.— Illustration how errors of 3 close-by roots scale with the distance between those roots. Upper dashed lines show analytic estimates of the errors one should expect when solving the three last roots of polynomial with a cubic solver - Equation (21), see also previous figures. Points with errorbars show the result from simulations when the three last roots were found numerically using the undivided polynomial, for cases when group of close-by roots was located at distances 1, 10^{-2} , 10^{-4} , and 10^{-6} from the origin of coordinate system. We see that when the characteristic distance from the origin of the last group of close-by roots is close to unity, $o \sim 1$, the accuracy of the numerical algorithm is well approximated by the analytical results. When the distance o is smaller, the accuracy of the numerical results scales with o^3 in the case when 3 roots are close.